arXiv:1607.03854v5 [cs.IT] 11 Apr 2017

The Partially Observable Hidden Markov Model and its Application to Keystroke Dynamics

John V. Monaco[a,*], Charles C. Tappert[b]

[a] *U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005, USA*
[b] *Pace University, Pleasantville, NY 10570, USA*

Abstract

The partially observable hidden Markov model is an extension of the hidden Markov Model in which the hidden state is conditioned on an independent Markov chain. This structure is motivated by the presence of discrete metadata, such as an event type, that may partially reveal the hidden state, but itself emanates from a separate process. Such a scenario is encountered in keystroke dynamics whereby a user's typing behavior is dependent on the text that is typed. Under the assumption that the user can be in either an active or passive state of typing, the keyboard key names are event types that partially reveal the hidden state due to the presence of relatively longer time intervals between words and sentences than between letters of a word. Using five public datasets, the proposed model is shown to consistently outperform other anomaly detectors, including the standard HMM, in biometric identification and verification tasks and is generally preferred over the HMM in a Monte Carlo goodness of fit test.

Keywords: hidden Markov model, keystroke biometrics, behavioral biometrics, time intervals, anomaly detection

1. Introduction

Time interval biometrics utilize the timestamps from a sequence of discrete events for the purpose of identification and verification. The sequence of time

*Corresponding author
Email address: john.v.monaco2.civ@mail.mil (John V. Monaco)
URL: http://www.vmonaco.com (John V. Monaco)

intervals, i.e., time between events, is of interest. Let t_n be the time of the n^{th} event. The sequence of time intervals is given by

$$\tau_n = t_n - t_{n-1} \ . \tag{1}$$

In the case an event has a duration, t_n marks the time of onset.

There are many scenarios in which timestamped events are labeled with some metadata, such as an event type. Consider a two-state model of keystroke dynamics in which the user can be in either an active or passive state of typing. Each new event, i.e., each key pressed, generates a new time interval. The keyboard key names can be considered event types that partially reveal the hidden state since the probability of being in either an active or passive state of typing may be greater depending on which key was pressed. Non-letter keys, such as punctuation and the Space key, indicate a greater probability of being in a passive state, as the typist often pauses between words and sentences as opposed to between letters in a word [1]. This reasoning extends to other activities, such as email, in which a user might be more likely to take an extended break after sending an email instead of receiving an email, and programming, in which a user may fix bugs quicker than making feature additions.

More generally, the sequence of emissions, i.e., the time intervals, comprise the system *behavior*, and the sequence of event types, i.e., the keys pressed, comprise the system *task*. While the time intervals reflect *how* the system behaves, the sequence of keys characterize *what* the system is doing. Keystroke dynamics aims to capture typing behavior by the time intervals between keystrokes, however this does not necessary include the key sequence itself which could emanate from a completely separate process. The key sequence may be assigned to the user for transcription, such as a passphrase for authentication, in which case the sequence has not been chosen at all by the user. Alternatively, in answering an open-ended question, even though the key sequence may reflect the user's writing style and linguistic behavior, it still pertains to the topic of the question. Keystroke dynamics aims to capture how the sequence was typed, not necessar-

ily what was typed. In both scenarios, the user's typing behavior depends on the key sequence, which itself may or may not reflect user behavior.

Motivated by the above scenario, we introduce the partially observable hidden Markov model (POHMM), an extension of the hidden Markov Model (HMM) in which the hidden state is conditioned on an independent Markov chain of event types. Despite an explosion in the number of model parameters, parameter estimation can still be performed in time that scales linearly with the number of observations. Marginal distributions of the model act as a fallback mechanism when novel event types are encountered during likelihood estimation. A parameter smoothing technique is used to regularize the parameter estimates and account for missing or infrequent data. Using five public datasets, the proposed model is shown to consistently outperform other anomaly detectors, including the standard HMM, in biometric identification and verification tasks and is generally preferred over the HMM in a Monte Carlo goodness of fit test.

The rest of this article is organized as follows. Section 2 briefly describes keystroke dynamics as a behavioral biometric. Section 3 introduces the POHMM, followed by a simulation study in Section 4 and a case study of the POHMM applied to keystroke dynamics in Section 5. Section 6 reviews previous modeling efforts for latent processes with partial observability and contains a discussion. Finally, Section 7 concludes the article. The POHMM is implemented in the `pohmm` Python package[1].

2. Keystroke dynamics

Keystrokes are typically detected by a *keylogger*, a program that records the sequence of keyboard keys pressed and released on a desktop or laptop computer. The resulting sequence of timings comprise a user's *keystroke dynamics*, wherein each keystroke contains a press timestamp and release timestamp. On commodity hardware, key event timestamps are commonly recorded with mil-

[1] Available at https://github.com/vmonaco/pohmm and through PyPI.

lisecond (ms) resolution, although the actual resolution is limited by the resolution of a global timer in the operating system kernel [2]. The global timer has a binning effect on the timestamps, as key press and release interrupts are handled only on the timer intervals. Despite this, the keystroke dynamics, which is thought to uniquely characterize an individual's typing behavior, can still be reliably detected, as the median time interval of the average typist (174 ms [1]) is about an order of magnitude greater than that of the timer. On Windows, the default timer updates every 15.625 ms (64 Hz) [3]; on Linux (X11) and Mac, a 10 ms (100 Hz) timer has traditionally been used [4]. The timer update rate varies greatly depending on system configuration and processing load.

As a behavioral biometric, keystroke dynamics enables low-cost and non-intrusive user identification and verification. Keystroke dynamics-based verification can be deployed remotely, often as a second factor to username-password verification. Some of the same attributes that make keystroke dynamics attractive as a behavioral biometric also present privacy concerns [5], as there exist numerous methods of detecting keystrokes without running software on the victim's computer. Recently, it has been demonstrated that keystrokes can be detected through a wide range of modalities including motion [6], acoustics [7], network traffic [8], and even WiFi signal distortion [9].

Keystroke time intervals emanate from a combination of physiology (e.g., age, gender, and handedness [10]), motor behavior (e.g., typing skill [1]), and higher-level cognitive processes [11]. A user's keystroke dynamics is also generally dependent on the typing task. For example, the time intervals observed during password entry are much different than those observed during email composition. For the purpose of biometric identification and verification, keystrokes are typically categorized by input type as follows[2].

Fixed-text: The keystrokes exactly follow a relatively short predefined sequence. The character sequence is known beforehand and entered without errors

[2]The keystrokes refer to the sequence of keys that were pressed, and the character sequence refers to the characters that ultimately appear on screen. Note that the keystrokes may vary to produce the same character sequence if typing corrections are allowed.

or corrections. This includes passwords, phone numbers, and personal identification numbers (PIN) without making corrections.

Constrained-text: The keystrokes roughly follow a predefined sequence. The character sequence is known beforehand, but typing errors and corrections are allowed. This includes case-insensitive passwords, passwords that were typed with corrections (i.e., including the Backspace key), or transcribed text. This type of input is used by some massively open online course (MOOC) providers in which the student must copy several sentences for the purpose of keystroke dynamics-based verification [12].

Free-text: The keystrokes do not follow a predefined sequence. The character sequence is unknown beforehand and usually consists of at least several sentences. The keystrokes collected as part of an essay question in an online test would be considered long free-text.

Due to the keyboard being one of the primary human-computer interfaces, it is also natural to consider keystroke dynamics as a modality for *continuous verification*, in which a verification decision is made upon each new key pressed [13]. Continuous verification holds the promise of greater security, as users are verified continuously throughout a session beyond the initial login, which is considered a form of *static verification*. Being a sequential model, the POHMM is straightforward to use for continuous verification, in addition to identification and static verification.

3. Partially observable hidden Markov model

The HMM is a finite-state model in which observed values at time t depend on an underlying latent process [14]. At the n^{th} time step t_n, a feature vector $\mathbf{x}_n$ is emitted and the system can be in any one of M hidden states, z_n. Let $\mathbf{x}_1^N$ be the sequence of observed emission vectors and z_1^N the sequence of hidden states, where N is the total number of observations. The basic HMM is defined by the recurrence relation,

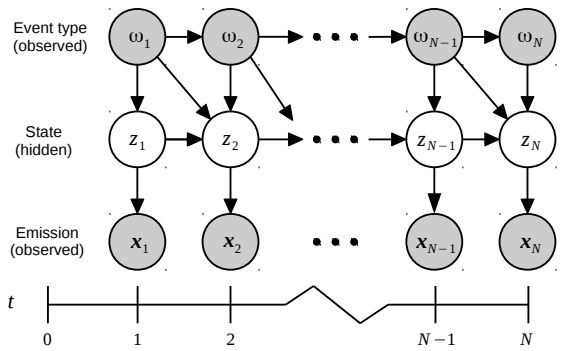

Figure 1: Partially observable hidden Markov model structure. Observed values (emission and event type) are shown in gray, hidden values (system state) are shown in white.

$$P\left(\mathbf{x}_1^{n+1}, z_1^{n+1}\right) = P\left(\mathbf{x}_1^n, z_1^n\right) P\left(\mathbf{x}_{n+1}|z_{n+1}\right) P\left(z_{n+1}|z_n\right) \ . \tag{2}$$

The POHMM is an extension of the HMM in which the hidden state and emission depend on an observed independent Markov chain. The POHMM is defined recursively by the formula,

$$P\left(\mathbf{x}_1^{n+1}, z_1^{n+1}\right) = P\left(\mathbf{x}_1^n, z_1^n\right) P\left(\mathbf{x}_{n+1}|z_{n+1}, \Omega_{n+1}\right) P\left(z_{n+1}|z_n, \Omega_n, \Omega_{n+1}\right) \tag{3}$$

where Ω_n and Ω_{n+1} are the observed event types at times t_n and t_{n+1}, respectfully, and Ω_1^N form an independent Markov chain. The POHMM structure is shown in Figure 1. The main differences from the HMM are that:

- The emission depends on both the unobserved hidden state and the observed event type.
- The hidden state depends on the previous and current observed event types in addition to the previous hidden state.

The event types come from a finite alphabet of size m. Thus, while the HMM has M hidden states, a POHMM with m event types has M hidden states per event type, for a total of $m \times M$ unique hidden states. The event type can be viewed as a partial indexing to a much larger state space. Each observed event type restricts the model to a particular subset of M hidden states, with differing

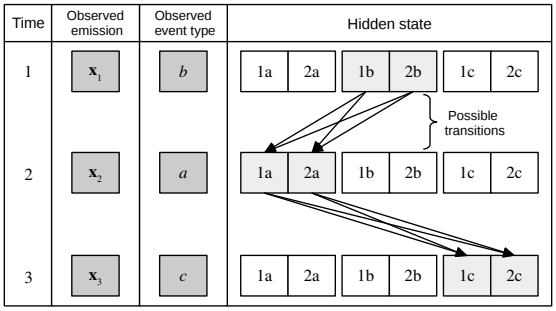

Figure 2: POHMM event types index a much larger state space. In this example, there are two hidden states and three event types. Given observed event type b at time 1, the system must be in one of the hidden states $\{1b, 2b\}$. The a observed at the next time step limits the possible transitions from $\{1b, 2b\}$ to $\{1a, 2a\}$.

probabilities of being in each hidden state, hence the partial observability. The POHMM starting and emission probabilities can be viewed as an HMM for each event type, and the POHMM transition probabilities as an HMM for each *pair* of event types.

To illustrate this concept, consider a POHMM with two hidden states and three event types, where $\Omega_1^3 = [b, a, c]$. At each time step, the observed event type limits the system to hidden states that have been conditioned on that event type, as demonstrated in Figure 2. Beginning at time 1, given observed event type $\Omega_1 = b$, the system must be in one of the hidden states $\{1b, 2b\}$. Event type $\Omega_2 = a$ observed at time 2 then restricts the possible transitions from $\{1b, 2b\}$ to $\{1a, 2a\}$. Generally, given any event type, the POHMM must be in one of M hidden states conditioned on that event type; therefore the state transitions, given the event types, have the same complexity as an HMM with an equal number of hidden states. Section 3.5 deals with situations where the event type is missing or has not been previously observed, in which case the marginal distributions (with the event type marginalized out) are used.

The POHMM parameters are defined similarly to the HMM, except all parameters are conditioned on the observed event type. Specifically, model parameters include $\pi[j|\omega]$, the probability of starting in state j given event type ω, and $a[i, j|\psi, \omega]$, the probability of transitioning from state i to state j, given event

types ψ and ω before and after the transition, respectfully[3]. Let $f(\cdot; \mathbf{b}[j|\omega])$ be the emission distribution that depends on hidden state j and event type ω, where $\mathbf{b}[j|\omega]$ parametrizes density function $f(\cdot)$. The complete set of parameters is denoted by $\theta = \{\pi, a, \mathbf{b}\}$, where a is the m^2M^2 transition matrix. While the total number of parameters in the HMM is $M + M^2 + MK$, where K is the number of free parameters in the emission distribution, the POHMM contains $mM + m^2M^2 + mMK$ parameters. After accounting for normalization constraints, the degrees of freedom (dof) is $m(M-1)+m^2M(M-1)+mMK$.

Marginal distributions, in which the event type is marginalized out, are also defined. Let $\pi[j]$ and $f(\cdot; \mathbf{b}[j])$ be the marginalized starting and emission probabilities, respectively. Similarly, the parameters $a[i, j|\omega]$, $a[i, j|\psi]$, and $a[i, j]$ are defined as the transition probabilities after marginalizing out the first, second, and both event types, respectively. *The POHMM marginal distributions are exactly equal to the corresponding HMM that ignores the event types.* This ensures that the POHMM is no worse than the HMM in case the event types provide little or no information as to the process being modeled. Computation of POHMM marginal distributions is covered in Section 3.5 and simulation results demonstrating this equivalence are in Section 4.

It may seem that POHMM parameter estimation (having all conditional parameters) becomes intractable, as the number of possible transitions between hidden states increases by a factor of m^2 over the HMM and all other parameters by a factor of m. In fact, all of the algorithms used for the POHMM are natural extensions of those used for the HMM. Parameter estimation remains linearly bounded in the number of observations, performed through a modification of the Baum-Welch (BW) algorithm. The convergence property of the modified BW algorithm is demonstrated analytically in Section 3.3 and empirically in Section 4. The rest of this section addresses the three main problems of the POHMM, taken analogously as the three main problems of the HMM:

[3]When a transition is involved, i and ψ always refer to the hidden state and event type, respectively, before the transition; j and ω refer to those after the transition.

1. Determine $P(\mathbf{x}_1^N|\Omega_1^N, \theta)$, the likelihood of an emission sequence given the model parameters and the observed event types.
2. Determine z_1^N, the maximum likelihood sequence of hidden states, given the emissions $\mathbf{x}_1^N$ and event types Ω_1^N.
3. Determine $\arg\max_{\theta\in\Theta} P(\mathbf{x}_1^N|\Omega_1^N, \theta)$, the maximum likelihood parameters θ for observed emission and event type sequences.

The first and third problems are necessary for identifying and verifying users in biometric applications, while the second problem is useful for understanding user behavior. The rest of this section reviews the solutions to each of these problems and other aspects of parameter estimation, including parameter initialization and smoothing.

3.1. Model likelihood

The likelihood of an emission sequence, given the model parameters θ and the observed event type sequence Ω_1^N, is denoted by $P(\mathbf{x}_1^N|\Omega_1^N)$[4]. In the HMM, $P(\mathbf{x}_1^N)$ can be computed efficiently by the forward procedure which defines a recurrence beginning at the start of the sequence. This procedure differs slightly for the POHMM due to the dependence on event types. Notably, the starting, transition, and emission parameters are all conditioned on the event type.

Let $\alpha_n[z_n, \Omega_n] \equiv P(\mathbf{x}_1^n, z_n|\Omega_n)$, i.e., the joint probability of emission sequence $\mathbf{x}_1^n$ and hidden state z_n, given event type Ω_n. Then, by the POHMM axiom (Equation 3), $\alpha_n[z_n, \Omega_n]$ can be computed recursively by the formula,

$$\begin{aligned} \alpha_{n+1}[z_{n+1}, \Omega_{n+1}] &= P(\mathbf{x}_{n+1}|z_{n+1}, \Omega_{n+1}) \\ &\quad \times \sum_{z_n} P(z_{n+1}|z_n, \Omega_n, \Omega_{n+1})\, \alpha_n[z_n, \Omega_n] \end{aligned} \tag{4}$$

$$\alpha_1[z_1, \Omega_1] = P(\mathbf{x}_1|z_1, \Omega_1)\, P(z_1|\Omega_1) \tag{5}$$

[4]For brevity, the dependence on θ is implied, writing $P(\mathbf{x}_1^N|\Omega_1^N, \theta)$ as $P(\mathbf{x}_1^N|\Omega_1^N)$.

where Equation 5 provides the initial condition. The modified forward algorithm is obtained by substituting the model parameters into Equations 4 and 5, where

$$\begin{aligned} \pi\left[j|\omega\right] &\equiv P\left(z_1 = j|\Omega_1 = \omega\right) && (6) \\ f(\mathbf{x}_n; \mathbf{b}\left[j,\omega\right]) &\equiv P\left(\mathbf{x}_n|z_n = j, \Omega_n = \omega\right) && (7) \\ a\left[i,j|\psi,\omega\right] &\equiv P\left(z_{n+1} = j|z_n = i, \Omega_n = \psi, \Omega_{n+1} = \omega\right) && (8) \end{aligned}$$

and $\alpha_n\left[j,\omega\right]$ is the sequence obtained after substituting the model parameters. The model likelihood is easily computed upon termination, since $P(\mathbf{x}_1^N|\Omega_1^N) = \sum_{j=1}^{M} \alpha_N\left[j,\omega\right]$ where $\omega = \Omega_N$.

A modified backward procedure is similarly defined through a backwards recurrence. Let $\beta_n\left[z_n,\Omega_n\right] \equiv P\left(\mathbf{x}_{n+1}^N|z_n,\Omega_n\right)$. Then under the POHMM axiom,

$$\begin{aligned} \beta_n\left[z_n,\Omega_n\right] &= \sum_{z_{n+1}} P\left(\mathbf{x}_{n+1}|z_{n+1},\Omega_{n+1}\right) && (9) \\ &\quad \times P\left(z_{n+1}|z_n,\Omega_n,\Omega_{n+1}\right)\beta_{n+1}\left[z_{n+1},\Omega_{n+1}\right] \\ \beta_N\left[z_N,\Omega_N\right] &= 1\,. && (10) \end{aligned}$$

where $\beta_n\left[j,\omega\right]$ is the sequence obtained after making the same substitutions.

Note that at each n, $\alpha_n\left[j,\omega\right]$ and $\beta_n\left[j,\omega\right]$ need only be computed for the observed $\omega = \Omega_n$, i.e., we don't care about event types $\omega \neq \Omega_n$. Therefore, only the hidden states (and not the event types) are enumerated in Equations 4 and 9 at each time step. Like the HMM, the modified forward and backward algorithms both take $O(M^2N)$ time to compute and can be stored in a $N \times M$ matrix.

3.2. Hidden state prediction

The maximum likelihood sequence of hidden states is efficiently computed using the event type-dependent forward and backward variables defined above. First, let the POHMM forward-backward variable $\gamma_n\left[z_n,\Omega_n\right] \equiv P\left(z_n|\Omega_n,\mathbf{x}_1^N\right)$, i.e., the posterior probability of hidden state z_n, given event type Ω_n and the

Algorithm 1 Modified Baum-Welch for POHMM parameter estimation.

1. **Initialization**
   Choose initial parameters θ_o and let $\theta \leftarrow \theta_o$.
2. **Expectation**
   Use θ, $\mathbf{x}_1^N$, Ω_1^N to compute $\alpha_n[j|\omega]$, $\beta_n[j|\omega]$, $\gamma_n[j,\omega]$, $\xi_n[i,j|\psi,\omega]$.
3. **Maximization**
   Update θ using the re-estimation formulae (Eqs. 13, 15, 16) to get $\dot{\theta} = \left\{\dot{\pi}, \dot{a}, \dot{\mathbf{b}}\right\}$.
4. **Regularization**
   Calculate marginal distributions and apply parameter smoothing formulae.
5. **Termination**
   If $\ln P\left(\mathbf{x}_1^N|\Omega_1^N,\dot{\theta}\right) - \ln P\left(\mathbf{x}_1^N|\Omega_1^N,\theta\right) < \epsilon$, stop; else let $\theta \leftarrow \dot{\theta}$ and go to step 2.

emission sequence $\mathbf{x}_1^N$. Let $\gamma_n[j,\omega]$ be the estimate obtained using the model parameters, making the same substitutions as above. Then $\gamma_n[j,\omega]$ is straightforward to compute using the forward and backward variables, given by

$$\begin{aligned}\gamma_n[j,\omega] &= \frac{\alpha_n[j|\omega]\,\beta_n[j|\omega]}{P(\mathbf{x}_1^N|\Omega_1^N)} \\ &= \frac{\alpha_n[j|\omega]\,\beta_n[j|\omega]}{\sum_{i=1}^{M}\alpha_n[i|\omega]\,\beta_n[i|\omega]}\end{aligned} \tag{11}$$

where $\omega = \Omega_n$. The sequence of maximum likelihood hidden states is taken as,

$$z_n = \arg\max_{1\le j\le M}\gamma_n[j,\omega]\ . \tag{12}$$

Similar to $\alpha_n[j|\omega]$ and $\beta_n[j|\omega]$, $\gamma_n[j,\omega]$ can be stored in a $N \times M$ matrix and takes $O\left(M^2N\right)$ time to compute. This is due to the fact that the event types are not enumerated at each step; the dependency on the event type propagates all the way to the re-estimated parameters, defined below.

3.3. Parameter estimation

Parameter estimation is performed iteratively, updating the starting, transition, and emission parameters using the current model parameters and observed sequences. In each iteration of the modified Baum-Welch algorithm, summarized in Algorithm 1, the model parameters are re-estimated using the POHMM forward, backward, and forward-backward variables. Parameters are set to initial

values before the first iteration, and convergence is reached upon a loglikelihood increase of less than ϵ.

3.3.1. Starting parameters

Using the modified forward-backward variable given by Equation 11, the POHMM starting probabilities are obtained directly by

$$\dot{\pi}\left[j|\omega\right] = \gamma_1\left[j|\omega\right] \tag{13}$$

where $\omega = \Omega_1$ and re-estimated parameters are denoted by a dot. Generally, it may not be possible to estimate $\dot{\pi}\left[j|\omega\right]$ for many ω due to there only being one Ω_1 (or several Ω_1 for multiple observation sequences). Parameter smoothing, introduced in Section 3.6, addresses the issue of missing and infrequent observations.

3.3.2. Transition parameters

In contrast to the HMM, which has M^2 transition probabilities, there are m^2M^2 unique transition probabilities in the POHMM. Let $\xi_n\left[z_n, z_{n+1}|\Omega_n, \Omega_{n+1}\right] \equiv P\left(z_{n+1}|z_n, \Omega_n, \Omega_{n+1}, \mathbf{x}_1^N\right)$, i.e., the probability of transitioning from state z_n to z_{n+1}, given event types Ω_n and Ω_{n+1} as well as the emission sequence. Substituting the forward and backward variable estimates based on model parameters, this becomes $\xi_n\left[i, j|\psi, \omega\right]$, given by

$$\xi_n\left[i, j|\psi, \omega\right] = \frac{\alpha_n\left[i, \omega\right] a\left[i, j|\psi, \omega\right] f\left(\mathbf{x}_{n+1}; \mathbf{b}\left[j|\omega\right]\right) \beta_n\left[j|\omega\right]}{P(\mathbf{x}_1^N|\Omega_1^N)} \,. \tag{14}$$

for $1 \leq n \leq N-1$, $\psi = \Omega_n$ and $\omega = \Omega_{n+1}$. The updated transition parameters are then calculated by

$$\dot{a}\left[i, j|\psi, \omega\right] = \frac{\sum_{n=1}^{N-1} \xi_n\left[i, j|\psi, \omega\right] \delta\left(\psi, \Omega_n\right) \delta\left(\omega, \Omega_{n+1}\right)}{\sum_{n=1}^{N-1} \gamma_n\left[i|\psi\right] \delta\left(\psi, \Omega_n\right) \delta\left(\omega, \Omega_{n+1}\right)} \tag{15}$$

where $\delta\left(\omega, \Omega_n\right) = 1$ if $\omega = \Omega_n$ and 0 otherwise. Note that $\dot{a}\left[i, j|\psi, \omega\right]$ depends only on the transitions between event types ψ and ω in Ω_1^N, i.e., where $\Omega_n = \psi$

and $\Omega_{n+1} = \omega$, as the summand in the numerator equals 0 otherwise. As a result, the updated transition probabilities can be computed in $O(M^2 N)$ time, the same as the HMM, despite there being $m^2 M^2$ unique transitions.

3.3.3. Emission parameters

For each hidden state and event type, the emission distribution parameters are re-estimated through the optimization problem,

$$\dot{\mathbf{b}}\left[j|\omega\right] = \arg\max_{\mathbf{b}\in\mathcal{B}} \sum_{n=1}^{N} \gamma_n\left[j|\omega\right] \ln f\left(\mathbf{x}_n; \mathbf{b}\right) \delta\left(\omega, \Omega_n\right) \ . \tag{16}$$

Closed-form expressions exist for a variety of emission distributions, including the log-normal. With the log-normal density defined as,

$$f(x; \eta, \rho) = \frac{1}{x\rho\sqrt{2\pi}} \exp\left[\frac{-\left(\ln x - \eta\right)^2}{2\rho^2}\right] \tag{17}$$

where η and ρ are the log-mean and log-standard deviation, respectively, emission parameter re-estimates are given by

$$\dot{\eta}\left[j|\omega\right] = \frac{\sum_{n=1}^{N} \gamma_n\left[j|\omega\right] \ln \tau_n \delta\left(\omega, \Omega_n\right)}{\sum_{n=1}^{N} \gamma_n\left[j|\omega\right] \delta\left(\psi, \Omega_n\right)} \tag{18}$$

and

$$\dot{\rho}^2\left[j|\omega\right] = \frac{\sum_{n=1}^{N} \gamma_n\left[j|\omega\right] \left(\ln \tau_n - \dot{\eta}_{j|\omega}\right)^2 \delta\left(\omega, \Omega_n\right)}{\sum_{n=1}^{N} \gamma_n\left[j|\omega\right] \delta\left(\psi, \Omega_n\right)} \tag{19}$$

for hidden state j, given event type ω. Note that the estimates for $\dot{\eta}\left[j|\omega\right]$ and $\dot{\rho}\left[j|\omega\right]$ depend only on the elements of $\gamma_n\left[j|\omega\right]$ where $\Omega_n = \omega$.

3.3.4. Convergence properties

The modified Baum-Welch algorithm for POHMM parameter estimation (Algorithm 1) relies on the principles of expectation maximization (EM) and is guaranteed to converge to a local maximum. The re-estimation formula (Equations 13, 15, and 16) are derived from inserting the model parameters from two successive iterations, θ and $\dot{\theta}$, into Baum's auxiliary function, $Q\left(\theta, \dot{\theta}\right)$,

and maximizing $Q\left(\theta, \dot{\theta}\right)$ with respect to the updated parameters. Convergence properties are evaluated empirically in Section 4, and Appendix B contains a proof of convergence, which follows that of the HMM.

3.4. Parameter initialization

Parameter estimation begins with parameter initialization, which plays an important role in the BW algorithm and may ultimately determine the quality of the estimated model, as EM guarantees only locally maximum likelihood estimates. This work uses an observation-based parameter initialization procedure that ensures reproducible parameter estimates, as opposed to random initialization. The starting and transition probabilities are simply initialized as

$$\pi\left[j|\omega\right] = \frac{1}{M} \tag{20}$$

$$a\left[i,j|\psi,\omega\right] = \frac{1}{M} \tag{21}$$

for all i, j, ψ, and ω. This reflects maximum entropy, i.e., uniform distribution, in the absence of any starting or transition priors.

Next, the emission distribution parameters are initialized. The strategy proposed here is to initialize parameters in such a way that there is a correspondence between hidden states from two different models. That is, for any two models with $M = 2$, hidden state $j = 1$ corresponds to the active state and $j = 2$ corresponds to the passive state. Using a log-normal emission distribution, this is accomplished by spreading the log-mean initial parameters. Let

$$\eta\left[\omega\right] = \frac{\sum_{n=1}^{N} \ln \mathbf{x}_n \delta\left(\omega, \Omega_n\right)}{\sum_{n=1}^{N} \delta\left(\omega, \Omega_n\right)} \tag{22}$$

and

$$\rho^2\left[\omega\right] = \frac{\sum_{n=1}^{N} (\ln \mathbf{x}_n - \eta\left[\omega\right])^2 \delta\left(\omega, \Omega_n\right)}{\sum_{n=1}^{N} \delta\left(\omega, \Omega_n\right)} \tag{23}$$

be the observed log-mean and log-variance for event type ω. The model param-

eters are then initialized as

$$\eta\left[j|\omega\right] = \eta\left[\omega\right] + \left(\frac{2h\left(j-1\right)}{M-1} - h\right)\rho\left[\omega\right] \tag{24}$$

and

$$\rho^2\left[j|\omega\right] = \rho^2\left[\omega\right] \tag{25}$$

for $1 \leq j \leq M$, where h is a bandwidth parameter. Using $h = 2$, initial states are spread over the interval $\left[\eta\left[\omega\right] - 2\rho\left[\omega\right], \eta\left[\omega\right] + 2\rho\left[\omega\right]\right]$, i.e., 2 log-standard deviations around the log-mean. This ensures that $j = 1$ corresponds to the state with the smaller log-mean, i.e., the active state.

3.5. Marginal distributions

When computing the likelihood of a novel sequence, it is possible that some event types were not encountered during parameter estimation. This situation arises when event types correspond to key names of freely-typed text and novel key sequences are observed during testing. A fallback mechanism (sometimes referred to as a "backoff" model) is typically employed to handle missing or sparse data, such as that used linguistics [15]. In order for the POHMM to handle missing or novel event types during likelihood calculation, the marginal distributions are used. This creates a two-level fallback hierarchy in which missing or novel event types fall back to the distribution in which the event type is marginalized out.

Let the probability of event type ω at time t_1 be $\pi\left[\omega\right]$, and the probability of transitioning from event type ψ to ω be denoted by $a\left[\psi, \omega\right]$. Both can be computed directly from the event type sequence Ω_1^N, which is assumed to be a first-order Markov chain. The marginal $\pi\left[j\right]$ is the probability of starting in hidden state j, in which the event type has been marginalized out,

$$\pi\left[j\right] = \sum_{\omega \in \Omega} \pi\left[j|\omega\right]\pi\left[\omega\right] \tag{26}$$

where Ω is the set of unique event types in Ω_1^N.

Marginal transition probabilities are also be defined. Let $a[i,j|\psi]$ be the probability of transitioning from hidden state i to hidden state j, given event type ψ while in hidden state i. The second event type for hidden state j has been marginalized out. This probability is given by

$$a[i,j|\psi] = \sum_{\omega\in\Omega} a[i,j|\psi,\omega]\, a[\psi,\omega] \ . \tag{27}$$

The marginal probability $a[i,j|\omega]$ is defined similarly by

$$a[i,j|\omega] = \frac{\sum_{\psi\in\Omega} a[i,j|\psi,\omega]\, a[\psi,\omega]}{\sum_{\psi\in\Omega} a[\psi,\omega]} \ . \tag{28}$$

Finally, the marginal $a[i,j]$ is the probability of transitioning from i to j,

$$a[i,j] = \frac{1}{m} \sum_{\psi\in\Omega} \sum_{\omega\in\Omega} a[i,j|\psi,\omega]\, a[\psi,\omega] \ . \tag{29}$$

No denominator is needed in Equation 27 since the normalization constraints of both transition matrices carry over to the left-hand side. Equation 29 is normalized by $\frac{1}{m}$ since $\sum_{\psi\in\Omega}\sum_{\omega\in\Omega} a[\psi,\omega] = m$.

The marginal emission distribution is a convex combination of the emission distributions conditioned on each of the event types. For normal and log-normal emissions, the marginal emission is simply a mixture of normals or log-normals, respectively. Let $\eta[j]$ and $\rho^2[j]$ be the log-mean and log-variance of the marginal distribution for hidden state j. The marginal log-mean is a weighted sum of the conditional distributions, given by

$$\eta[j] = \sum_{\omega\in\Omega} \Pi[\omega]\, \mu[j|\omega] \tag{30}$$

where $\Pi[\omega]$ is the stationary probability of event type ω. This can be calculated directly from the event type sequence Ω_1^N,

$$\Pi[\omega] = \frac{1}{N} \sum_{n=1}^{N} \delta(\omega, \Omega_n) \ . \tag{31}$$

Similarly, the marginal log-variance is a mixture of log-normals given by

$$\rho^2[j] = \sum_{\omega \in \Omega} \Pi[\omega] \left[(\eta[j|\omega] - \eta[j])^2 + \rho^2[j|\omega] \right] . \tag{32}$$

Marginalized distribution parameters for normal emission is exactly the same.

3.6. Parameter smoothing

While the marginal distributions can be used to handle missing or novel data during likelihood calculation, parameter smoothing handles missing or infrequent data during parameter estimation. The purpose of parameter smoothing is twofold. First, it acts as a kind of regularization to avoid overfitting, a problem often encountered when there is a large number of parameters and small number of observations. Second, parameter smoothing provides superior estimates in case of missing or infrequent data.

For motivation, consider a keystroke sequence of length N. Including English letters and the Space key, there are at most 27 unique keys and 729 unique digrams (subsequences of length 2). Most of these will rarely, or never, be observed in a sequence of English text. Parameter smoothing addresses this issue by re-estimating the parameters that depend on low-frequency observations using a mixture of the marginal distribution. The effect is to bias parameters that depend on event types with low frequency toward the marginals, for which there exist more observations and higher confidence, while parameters that depend on event types with high frequency will remain unchanged.

Smoothing weights for the starting and emission parameters are defined as

$$w_\omega = 1 - \frac{1}{1 + f(\omega)} \tag{33}$$

where $f(\omega) = \sum_{t=1}^{N} \delta(\omega, \Omega_n)$ is the frequency of event type ω in the sequence Ω_1^N. The POHMM starting probabilities are then smoothed by

$$\tilde{\pi}[j|\omega] = w_\omega \pi[j|\omega] + (1 - w_\omega) \pi[j] \tag{34}$$

with smoothed parameter estimates denoted by a tilde, and emission parameters are smoothed by

$$\tilde{\mathbf{b}}\left[j|\omega\right] = w_{\omega}\mathbf{b}\left[j|\omega\right] + \left(1 - w_{\omega}\right)\mathbf{b}\left[j\right] \ . \tag{35}$$

As N increases, event type frequencies increase and the effect of parameter smoothing is diminished, while parameters conditioned on infrequent or missing event types are biased toward the marginal. This ensures that the conditional parameters remain asymptotically unbiased as $N \to \infty$.

The smoothing weights for transition probabilities follow similar formulae. Let $f\left(\psi,\omega\right) = \sum_{t=1}^{N-1}\delta\left(\psi,\Omega_n\right)\delta\left(\omega,\Omega_{n+1}\right)$, i.e., the frequency of event type ψ followed by ω in the sequence Ω_1^N. Weights for the conditional and marginal transition probabilities are defined as

$$\begin{aligned} w_{\psi} &= \frac{1}{f\left(\psi,\omega\right) + f\left(\omega\right)} \\ w_{\omega} &= \frac{1}{f\left(\psi,\omega\right) + f\left(\psi\right)} \\ w_{\psi,\omega} &= 1 - \left(w_{\psi} + w_{\omega}\right) \\ w &= 0 \end{aligned} \tag{36}$$

where $w_{\psi,\omega} + w_{\psi} + w_{\omega} + w = 1$. The smoothed transition matrix is given by

$$\tilde{a}\left[i,j|\psi,\omega\right] = w_{\psi,\omega}a\left[i,j|\psi,\omega\right] + w_{\psi}a\left[i,j|\psi\right] + w_{\omega}a\left[i,j|\omega\right] + wa\left[i,j\right] \ . \tag{37}$$

In this strategy, the weight for the marginal $a\left[i,j\right]$ is 0, although in other weighting schemes, w could be non-zero.

4. Simulation study

It is important for statistical models and their implementations to be consistent. This requires that parameter estimation be both convergent and asymptotically unbiased. While POHMM consistency is theoretically guaranteed (proof in Appendix B), consistency of the POHMM implementation under several different scenarios is validated in this section using computational methods. First,

a model is initialized with parameters θ_o. From this model, S samples are generated, each containing N time intervals. For each sample, the best-estimate parameters $\hat{\theta}$ are computed using the modified BW algorithm (Algorithm 1). Let $\hat{\theta}_N$ be the parameters determined by the modified BW algorithm for an observed sequence of length N generated from a POHMM with true parameters θ_o. Consistency requires that

$$\lim_{N\to\infty} \frac{|\hat{\theta}_N - \theta_o|}{\max_{\hat{\theta}} |\hat{\theta}_N - \theta_o|} = \mathbf{0} \tag{38}$$

insensitive to the choice of θ_o. As N increases, parameter estimation should be able to recover the true model parameters from the observed data. Four different scenarios are considered:

1. Train a POHMM (without smoothing) on POHMM-generated data.
2. Train a POHMM (with smoothing) on POHMM-generated data.
3. Train a POHMM (without smoothing) using emissions generated from an HMM and random event types.
4. Train an HMM using emissions from a POHMM (ignore event types).

Convergence is theoretically guaranteed for scenarios 1 and 2. The first scenario tests the POHMM implementation without parameter smoothing and should yield unbiased estimates. Scenario 2 evaluates the POHMM implementation with parameter smoothing, whose effect diminishes as N increases. Consequently, the smoothed POHMM estimates approach that of the unsmoothed POHMM, and results should also indicate consistency.

Scenario 3 is a POHMM trained on an HMM, and scenario 4 is an HMM trained on a POHMM. In scenario 3, the underlying process is an HMM with the same number of hidden states as the POHMM, and the observed event types are completely decorrelated from the HMM. As a result, the event types do *not* partially reveal the hidden state. In this case, the POHMM marginal distributions, in which the event type is marginalized out, should converge to the HMM. Finally, scenario 4 simply demonstrates the inability of the HMM

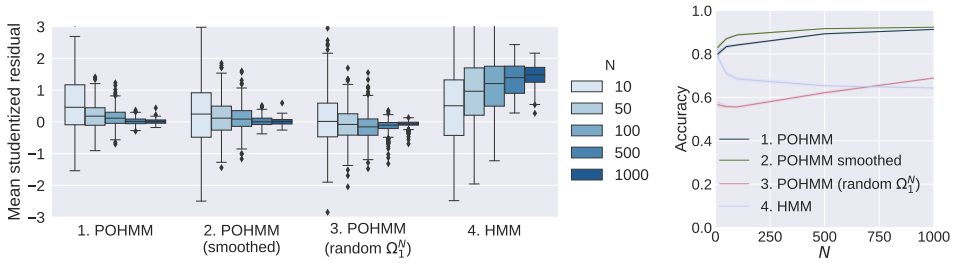

(a) Studentized emission residuals.

(b) Hidden state classification accuracy.

Figure 3: Simulation study results. In 1 and 2, a POHMM is trained on data generated from a POHMM; in 3, a POHMM is trained on data generated from an HMM (using random event types); in 4, an HMM is trained on data generated from a POHMM (ignoring event types).

to capture the dependence on event types, and results should indicate biased estimates.

For scenarios 1, 2 and 4, a POHMM with 3 event types and 2 hidden states is initialized to generate the training data. The emission distribution is a univariate Gaussian with parameters chosen to be comparable to human key-press time intervals, and transition probabilities follow a uniform distribution. The emission and event type sequences are sampled from the POHMM and used to fit the model. In scenario 3, an HMM generates the emission sequence $\mathbf{x}_1^N$, and the event type sequence Ω_1^N is chosen randomly from the set of 3 event types, reflecting no dependence on event types. In this case, only the POHMM marginal distribution parameter residuals are evaluated, as these should approximate the underlying HMM. For each value of N in each scenario, 400 length-N samples are generated and used to train the corresponding model.

Figure 3a contains the mean studentized residuals for emission parameters of each model, and Figure 3b shows the hidden state classification accuracies (where chance accuracy is $\frac{1}{2^N}$). Both the unsmoothed and smoothed POHMM residuals tend toward 0 as N increases, indicating consistency. The marginal residuals for the POHMM with random event types also appear unbiased, an indication that the POHMM marginals, in which the event type is marginalized

Table 1: Keystroke dataset summary. Columns 4-7 indicate: number of users, samples per user, keystrokes per sample, and $\bar{\tau}$=mean press-press latency (ms).

Dataset	Source	Category	Users	Samples/user	Keys/sample	$\bar{\tau}$ (ms)
Password	[16]	Short fixed	51	400	11	249
Keypad	[17]	Short fixed	30	20	11	376
Mobile	[18]	Short fixed	51	20	11	366
Fable	[19]	Long constrained	60	4	100	264
Essay	[20]	Long free	55	6	500	284

out, are asymptotically equivalent to the HMM. Finally, the HMM residuals, when trained on data generated from a POHMM, appear biased as expected when the event types are ignored. Similar results in all scenarios are seen for the transition probability residuals (not shown), and we confirmed that these results are insensitive to the choice of θ_o.

5. Case study: keystroke dynamics

Five publicly-available keystroke datasets are analyzed in this work, summarized in Table 1. The *password*, *keypad*, and *mobile* datasets contain short fixed-text input in which all the users in each dataset typed the same 10-character string followed by the Enter key: ".tie5Roanl" for the password dataset [16] and "9141937761" for the keypad [17] and mobile datasets [18]. Samples that contained errors or more than 11 keystrokes were discarded. The password dataset was collected on a laptop keyboard equipped with a high-resolution clock (estimated accuracy to within ± 200 μs [2]), while the timestamps in all other datasets were recorded with millisecond resolution (see discussion in Section 2 on timestamp resolution). The keypad dataset used only the 10-digit numeric keypad located on the right side a standard desktop keyboard, and the mobile dataset used an Android touchscreen keypad with similar layout. In addition to timestamps, the mobile dataset contains accelerometer, gyroscope, screen location, and pressure sensor features measured on each key press and release.

The *fable* dataset contains long constrained-text input from 60 users who each copied 4 different fables or nursery rhymes [19, 20]. Since mistakes were

permitted, the keystrokes for each copy task varied, unlike the short fixed-text datasets above. The *essay* dataset contains long free-text input from 55 users who each answered 6 essay-style questions as part of a class exercise [20]. Both the fable and essay datasets were collected on standard desktop and laptop keyboards. For this work, the fable samples were truncated to each contain exactly 100 keystrokes and the essay samples to each contain exactly 500 keystrokes.

Each keystroke event contains two timing features,

$$\tau_n = t_n^P - t_{n-1}^P \tag{39}$$

$$d_n = t_n^R - t_n^P \tag{40}$$

where t_n^P and t_n^R are the press and release timestamps of the n^{th} keystroke, respectively; τ_n is the press-press time interval and d_n is the key-hold duration. Note that other timing features, such as release-release and release-press intervals, can be calculated by a linear combination of the above two features.

Each user's keystroke dynamics are modeled by a POHMM with log-normal emission and two hidden states, all conditioned on the keyboard keys as the observed event types. A two-state model is the simplest model of non-homogeneous behavior, as one state implies a sequence of independent and identically distributed (i.i.d.) observations. The two hidden states correspond to the active and passive states of the user, in which relatively longer time intervals are observed in the passive state. Given the hidden state and the observed event type, the keystroke time intervals τ_n and d_n are each modeled by a log-normal distribution (Equation 17), where $\eta[j|\omega]$ and $\rho[j|\omega]$ are the log-mean and log-standard deviation, respectively, in hidden state j given observed key ω.

The POHMM parameters are determined using Algorithm 1, and convergence is achieved after a loglikelihood increase less than 10^{-6} or 1000 iterations, whichever is reached first. As an example, the marginal key-press time interval distributions for each hidden state are shown in Figure 4 for two randomly selected samples. The passive state in the free-text model has a heavier tail than the fixed-text, while the active state distributions in both models are compa-

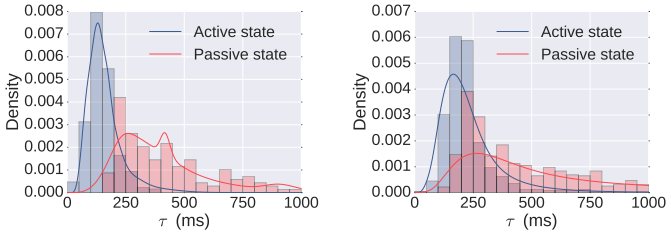

(a) Fixed-text (b) Free-text

Figure 4: POHMM marginal distributions showing a separation between active and passive typing states. The marginal distributions are mixtures of log-normals conditioned on the key names. Histograms show the empirical time interval distributions in each hidden state.

rable. The rest of this section presents experimental results for a goodness of fit test, identification, verification, and continuous verification. Source code to reproduce the experiments in this article is available[5].

5.1. Goodness of fit

To determine whether the POHMM is consistent with observed data, a Monte Carlo goodness of fit test is performed. The test proceeds as follows. For each keystroke sample (using the key-press time intervals only), the model parameters $\hat{\theta}_m$ are determined. The area test statistic between the model and empirical distribution is then taken. The area test statistic is a compromise between the Kolmogorov-Smirnov (KS) test and Cramér-von Mises test [21],

$$A = \int \left| P_D(\tau) - P_M\left(\tau|\hat{\theta}_m\right) \right| \mathrm{d}\tau \tag{41}$$

where P_D is the empirical cumulative distribution and P_M is the model cumulative distribution. The POHMM marginal emission density is given by

$$g\left(\mathbf{x};\theta\right) = \sum_{\omega\in\Omega}\sum_{j=1}^{M} \Pi\left[\omega\right]\Pi\left[j\right] f\left(\mathbf{x};\mathbf{b}\left[j|\omega\right]\right) \tag{42}$$

[5]Code to reproduce experiments: https://github.com/vmonaco/pohmm-keystroke

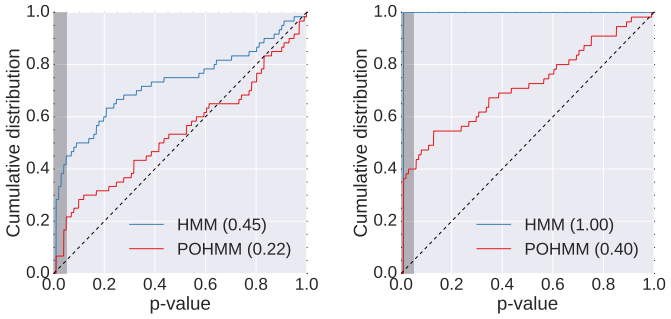

(a) Constrained-text (b) Free-text

Figure 5: Keystroke goodness of fit p-value distributions testing the null hypothesis that the model is consistent with the data. Proportions of rejected samples at the 0.05 significance level are shown in parentheses. If the null hypothesis was true, i.e., the model was actually consistent with the keystroke data, then p-values would follow a uniform distribution shown by the dashed black line.

where $\Pi[j]$ is the stationary probability of hidden state j and $\Pi[\omega]$ is the stationary probability of event type ω. Using the fitted model with parameters $\hat{\theta}_m$, a surrogate data sample the same size as the empirical sample is generated. Estimated parameters $\hat{\theta}_s$ are determined using the surrogate sample in a similar fashion as the empirical sample. The area test statistic between the surrogate-data-trained model and surrogate data is computed, given by A_s. This process repeats until enough surrogate statistics have accumulated to reliably determine $P(|A_s - \langle A_s \rangle| > |A - \langle A_s \rangle|)$. The biased p-value is given by

$$\frac{I\left(|A_s - \langle A_s \rangle| > |A - \langle A_s \rangle|\right) + 1}{S + 1} \tag{43}$$

where $I(\cdot)$ is the indicator function. Testing the null hypothesis, that the model is consistent with the data, requires fitting $S + 1$ models (1 empirical and S surrogate samples).

The test is performed for both the HMM and the POHMM for each user in the fable and essay datasets, using the key-press time intervals only. The resulting p-value distributions are shown in Figure 5. The shaded area represents a 0.05 significance level in which the null hypothesis is rejected. In the fable dataset, the HMM is rejected for 45% of users, while the POHMM is rejected for

22% of users. The HMM is rejected for 100% of users in the essay dataset, and the POHMM is rejected for 40% of users. If the POHMM truly reflected typing behavior (i.e., the null hypothesis was actually true), the p-values would follow a uniform distribution shown by the dashed black line. In both experiments, the POHMM is largely preferred over the HMM.

5.2. Identification and verification

The POHMM can be used for both identification and verification. Given a query sample from an unknown user, identification is performed by choosing the maximum likelihood model from a population of models. Verification is performed by comparing the claimed user model's loglikelihood to a threshold.

Identification and verification results are obtained for each keystroke dataset and four benchmark anomaly detectors in addition to the POHMM. The password dataset uses a validation procedure similar to Killourhy and Maxion [16], except only samples from the 4th session (repetitions 150-200) are used for training and sessions 5-8 (repetitions 201-400) for testing. For the other datasets, results are obtained through a stratified cross-fold validation procedure with the number of folds equal to the number of samples per user: 20 for keypad and mobile, 4 for fable, and 6 for essay. In each fold, one sample from each user is retained as a query and the remaining samples are used for training.

Identification accuracy (ACC) is measured by the proportion of correctly classified query samples. Verification performance is measured by the user-dependent equal error rate (EER), the point on the receiver operating characteristic (ROC) curve at which the false rejection rate (FRR) and false acceptance rate (FAR) are equal. Each query sample is compared against every model in the population, only one of which will be genuine. The resulting loglikelihood is normalized using the minimum and maximum loglikelihoods from every model in the population to obtain a normalized score between 0 and 1. Confidence intervals for both the ACC and EER are obtained over users in each dataset, as in [16].

Table 2: Identification accuracy rates. Bold indicates systems that are not significantly worse than the best system. Mobile+ includes mobile sensor features in addition to time intervals.

	Manhattan	Manhattan (Scaled)	SVM (One-class)	HMM	POHMM
Password	0.510 (0.307)	0.662 (0.282)	0.465 (0.293)	0.467 (0.295)	**0.789 (0.209)**
Keypad	**0.623 (0.256)**	**0.713 (0.200)**	0.500 (0.293)	0.478 (0.287)	**0.748 (0.151)**
Mobile	0.290 (0.230)	0.528 (0.237)	0.267 (0.229)	0.303 (0.265)	**0.607 (0.189)**
Mobile+	0.647 (0.250)	**0.947 (0.104)**	0.857 (0.232)	0.937 (0.085)	**0.971 (0.039)**
Fable	0.492 (0.332)	0.613 (0.314)	0.571 (0.235)	0.392 (0.355)	**0.887 (0.175)**
Essay	0.730 (0.320)	**0.839 (0.242)**	0.342 (0.302)	0.303 (0.351)	**0.909 (0.128)**

Benchmark anomaly detectors include Manhattan distance, scaled Manhattan distance, one-class support vector machine (SVM), and a two-state HMM. The Manhattan, scaled Manhattan, and one-class SVM operate on fixed-length feature vectors, unlike the HMM and POHMM. Timing feature vectors for the password, keypad, and mobile datasets are formed by the 11 press-press latencies and 10 durations of each 11-keystroke sample for a total of 21 timing features. The mobile sensors provide an additional 10 features for each keystroke event for a total of 131 features. For each event, the sensor features include: acceleration (meters/second^2) and rotation (radians/second) along three orthogonal axes (6 features), screen coordinates (2 features), pressure (1 feature), and the length of the major axis of an ellipse fit to the pointing device (1 feature). Feature vectors for the fable and essay datasets are each comprised of a set of 218 descriptive statistics for various keystroke timings. Such timing features include the sample mean and standard deviation of various sets of key durations, e.g., consonants, and latency between sets of keys, e.g., from consonants to vowels. For a complete list of features see [19, 22]. The feature extraction also includes a rigorous outlier removal step that excludes observations outside a specified confidence interval. A hierarchical fallback scheme accounts for missing or infrequent observations.

The Manhattan anomaly detector uses the negative Manhattan distance to the mean template vector as a confidence score. For the scaled Manhattan detector, features are first scaled by the mean absolute deviation over the entire dataset. This differs slightly from the scaled Manhattan in [16], which uses the mean absolute deviation of each user template. The global (over the entire

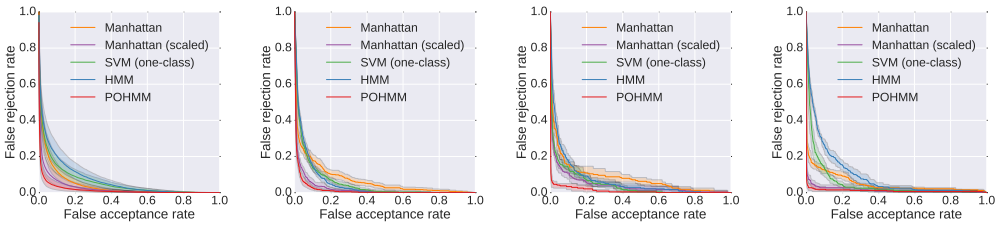

(a) Password (b) Keypad (c) Fixed-text (d) Free-text

Figure 6: Keystroke ROC curves. Bands show the 95% confidence intervals.

Table 3: User-dependent EER. Bold indicates systems that are not significantly worse than the best system. Mobile+ includes mobile sensor features in addition to time intervals.

	Manhattan	Manhattan (scaled)	SVM (one-class)	HMM	POHMM
Password	0.088 (0.069)	0.062 (0.064)	0.112 (0.088)	0.126 (0.099)	**0.042 (0.051)**
Keypad	0.092 (0.069)	**0.053 (0.030)**	0.110 (0.054)	0.099 (0.050)	**0.053 (0.025)**
Mobile	0.194 (0.101)	**0.097 (0.057)**	0.170 (0.092)	0.168 (0.085)	**0.090 (0.054)**
Mobile+	0.084 (0.061)	**0.009 (0.027)**	**0.014 (0.033)**	0.013 (0.021)	**0.006 (0.014)**
Fable	0.085 (0.091)	0.049 (0.060)	0.099 (0.106)	0.105 (0.092)	**0.031 (0.077)**
Essay	0.061 (0.092)	0.028 (0.052)	0.098 (0.091)	0.145 (0.107)	**0.020 (0.046)**

dataset) mean absolute deviation is used in this work due to the low number of samples per user in some datasets. The one-class SVM uses a radial basis function (RBF) kernel and 0.5 tolerance of training errors, i.e., half the samples will become support vectors. The HMM is exactly the same as the POHMM (two hidden states and log-normal emissions), except event types are ignored.

Identification and verification results are shown in Tables 2 and 3, respectively, and ROC curves are shown in Figure 6. The best-performing anomaly detectors in Tables 2 and 3 are shown in bold. The set of best-performing detectors contains those that are not significantly worse than the POHMM, which achieves the highest performance in every experiment. The Wilcoxon signed-rank test is used to determine whether a detector is significantly worse than the best detector, testing the null hypothesis that a detector has the same performance as the POHMM. A Bonferroni correction is applied to control the family-wise error rate, i.e., the probability of falsely rejecting a detector that is

actually in the set of best-performing detectors [23]. At a 0.05 significance level, the null hypothesis is rejected with a p-value not greater than $\frac{0.05}{4}$ since four tests are applied in each row. The POHMM achieves the highest identification accuracy and lowest equal error rate for each dataset. For 3 out of 6 datasets in both sets of experiments, all other detectors are found to be significantly worse than the POHMM.

5.3. Continuous verification

Keystrokes are continuously generated as a user interacts with a computer system, and in an intrusion detection scenario it is desirable to detect an impostor within as few keystrokes as possible. This differs from the static verification scenario in the previous section in which verification performance is evaluated over an entire session. Instead, continuous verification requires a verification decision to be made upon each new keystroke [13].

Continuous verification is enforced through a penalty function in which each new keystroke incurs a non-negative penalty within a sliding window. The penalty at any given time can be thought of as the inverse of trust. As behavior becomes more consistent with the model, the cumulative penalty within the window can decrease, and as it becomes more dissimilar, the penalty increases. The user is rejected if the cumulative penalty within the sliding window exceeds a threshold. The threshold is chosen for each sample such that the genuine user is never rejected, analogous to a 0% FRR in static verification. An alternative to the penalty function is the penalty-and-reward function in which keystrokes incur either a penalty or a reward (i.e., a negative penalty) [24]. In this work, the sliding window replaces the reward since penalties outside the window do not contribute towards the cumulative penalty.

The penalty of each new event is determined as follows. The marginal probability of each new event, given the preceding events, is obtained from the forward lattice, α, given by

$$P\left(\mathbf{x}_{n+1}|\mathbf{x}_1^n\right) = P\left(\mathbf{x}_1^{n+1}\right) - P\left(\mathbf{x}_1^n\right) \tag{44}$$

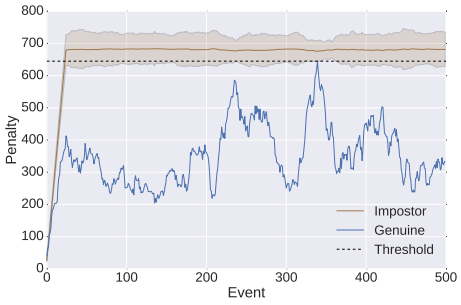

Figure 7: Continuous verification example. Bands show the 95% confidence interval. In this example, impostors are detected after an average of 81 keystrokes.

When a new event is observed, the likelihood is obtained under every model in a population of U models. The likelihoods are ranked, with the highest model given a rank of 0, and the lowest a rank of $U-1$. The rank of the claimed user's model is the incurred penalty. Thus, if a single event is correctly matched to the genuine user's model, a penalty of 0 is incurred; if it scores the second highest likelihood, a penalty of 1 is incurred, etc. The rank penalty is added to the cumulative penalty in the sliding window, while penalties outside the window are discarded. A window of length 25 is used in this work.

Continuous verification performance is reported as the number of events (up to the sample length) that can occur before an impostor is detected. This is determined by increasing the penalty threshold until the genuine user is never rejected by the system. Since the genuine user's penalty is always below the threshold, this is the maximum number of events that an impostor can execute before being rejected by the system while the genuine user is never rejected.

An example of the penalty function for genuine and impostor users is shown in Figure 7. The decision threshold is set to the maximum penalty incurred by the genuine user so that a false rejection does not occur. The average penalty for impostor users with 95% confidence interval is shown. In this example, the impostor penalties exceed the decision threshold after 81 keystrokes on average. Note that this is different than the average imposter penalty, which exceeds the threshold after 23 keystrokes.

Table 4: Continuous verification average maximum rejection time: the number of events that occur before an impostor is detected given the genuine user is not falsely rejected.

	HMM	POHMM
Password	5.64 (2.04)	3.42 (2.04)
Keypad	4.54 (2.09)	3.45 (1.73)
Mobile	5.63 (2.18)	4.29 (2.02)
Mobile+	0.15 (0.65)	0.12 (0.57)
Fable	33.63 (15.47)	20.81 (9.07)
Essay	129.36 (95.45)	55.18 (68.31)

For each dataset, the average maximum rejection time (AMRT) is determined, shown in Table 4. The maximum rejection time (MRT) is the maximum number of keystrokes needed to detect an impostor without rejecting the genuine user. The MRT is determined for each combination of impostor query sample and user model in the dataset to get the AMRT. The POHMM has a lower AMRT than the HMM for every dataset, and less than half that of the HMM for free-text input.

6. Discussion

There have been several generalizations of the standard HMM to deal with hidden states that are partially observable in some way. These models are referred to as partly-HMM [25], partially-HMM [26], and context-HMM [27].

The partly-HMM is a second order model in which the first state is hidden and the second state is observable [25]. In the partly-HMM, both the hidden state and emission at time t_n depend on the observation at time t_{n-1}. The partly-HMM can be applied to problems that have a transient underlying process, such as gesture and speech recognition, as opposed to a piecewise stationary process that the HMM assumes [28]. Parameter estimation is performed by the EM algorithm, similar to the HMM.

Partially observable states can also come in the form of partial and uncertain ground truth regarding the hidden state at each time step. The partially-HMM

addresses this scenario, in which an uncertain hidden state label may be observed at each time step [26]. The probability of observing the uncertain label and the probability of the label being correct, were the true hidden state known, are controlled by parameters p_{obs} and p_{true}, respectively. Thus, the probability of observing a correct label is $p_{obs} \times p_{true}$. This model is motivated by language modeling applications in which manually labeling data is expensive and time consuming. Similar to the HMM, the EM algorithm can be used for estimating the parameters of the partially-HMM [26].

Past observations can also provide context for the transition and emission probabilities in an HMM. Forchhammer and Rissanen [27] proposed the context-HMM, in which the transition and emission probabilities at time t_n are conditioned on a function of the previous observations up to time t_{n-1}. The context-HMM has information theoretic motivations, with applications such as image compression [29]. Used in this way, the neighboring pixels in an image can provide context for the emission and transition probabilities.

There are two scenarios in which previous models of partial observability fall short. The first is when there is missing data during parameter estimation, such missing context, and the second is when there is missing or novel data during likelihood calculation. A possible solution to these problems uses the explicit marginal emission and transition distributions, where, e.g., the context is marginalized out. While none of the above models possess this property, the POHMM, described in Section 3, has explicit marginal distributions that are used when missing or novel data are encountered. Additionally, parameter smoothing uses the marginal distributions to regularize the model and improve parameter estimates.

The POHMM is different from the partly-HMM [25], being a first order model, and different from the partially-HMM [26], since it doesn't assume a partial labeling. The POHMM is most similar to the context-HMM [27] in the sense that emission and transition probabilities are conditioned on some observed values. Despite this, there are several important differences between the POHMM and context-HMM.

The hidden state of the context-HMM at time t_n is dependent on the emissions up to time t_{n-1}, whereas the POHMM hidden state at time t_n depends on observations from a separate process at times t_n and t_{n-1}. This enables the POHMM to characterize system behavior that depends on an independent Markov chain, which emanates from a completely separate process. Such a scenario is encountered in keystroke dynamics, whereby typing behavior depends on the text that is being typed, but the text itself is not considered part of the keystroke dynamics. This distinction is not made in the context-HMM, as the context is based on the previously-observed emissions. Additionally, the context-HMM, as original described, contains only discrete distributions and lacks explicit marginal distributions; therefore it is unable to account for missing or novel data during likelihood calculation, as would be needed in free-text keystroke dynamics.

7. Conclusions

This work introduced the POHMM, an extension of the HMM in which the hidden states are partially observable through an independent Markov chain. Computational complexities of POHMM parameter estimation and likelihood calculation are comparable to that of the HMM, which are linear in the number of observations. POHMM parameter estimation also inherits the desirable properties of expectation maximization, as a modified Baum-Welch algorithm is employed. A case study of the POHMM applied to keystroke dynamics demonstrates superiority over leading alternative models on a variety of tasks, including identification, verification, and continuous verification. Given the ubiquity of timestamped events, other modalities of temporal behavior are of interest, and future work will explore the range of POHMM applications that extend beyond keystroke dynamics.

References

[1] T. A. Salthouse, Perceptual, cognitive, and motoric aspects of transcription typing., Psychological bulletin 99 (3) (1986) 303.

[2] K. Killourhy, R. Maxion, The effect of clock resolution on keystroke dynamics, in: Recent Advances in Intrusion Detection, Springer, 331–350, 2008.

[3] Microsoft, Timers, Timer Resolution, and Development of Efficient Code, URL `http://download.microsoft.com/download/3/0/2/3027D574-C433-412A-A8B6-5E0A75D5B237/Timer-Resolution.docx`, visited on 2016-12-28, 2010-06-16.

[4] Y. Etsion, D. Tsafrir, D. G. Feitelson, Effects of clock resolution on the scheduling of interactive and soft real-time processes, in: ACM SIGMETRICS Performance Evaluation Review, vol. 31, ACM, 172–183, 2003.

[5] J. V. Monaco, C. C. Tappert, Obfuscating Keystroke Time Intervals to Avoid Identification and Impersonation, arXiv preprint arXiv:1609.07612 .

[6] H. Wang, T. Tsung-Te Lai, R. R. Choudhury, MoLe: Motion Leaks through Smartwatch Sensors, in: The 21st Annual International Conference on Mobile Computing and Networking, ACM, 2015.

[7] D. Asonov, R. Agrawal, Keyboard acoustic emanations, in: Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on, IEEE, 3–11, 2004.

[8] D. X. Song, D. Wagner, X. Tian, Timing Analysis of Keystrokes and Timing Attacks on SSH., in: USENIX Security Symposium, vol. 2001, 2001.

[9] K. Ali, A. X. Liu, W. Wang, M. Shahzad, Keystroke recognition using wifi signals, in: Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, ACM, 90–102, 2015.

[10] S. Z. S. Idrus, E. Cherrier, C. Rosenberger, P. Bours, Soft biometrics for keystroke dynamics: Profiling individuals while typing passwords, Computers & Security 45 (2014) 147–155.

[11] D. G. Brizan, A. Goodkind, P. Koch, K. Balagani, V. V. Phoha, A. Rosenberg, Utilizing linguistically-enhanced keystroke dynamics to predict typist cognition and demographics, International Journal of Human-Computer Studies .

[12] A. Maas, C. Heather, C. T. Do, R. Brandman, D. Koller, A. Ng, Offering Verified Credentials in Massive Open Online Courses: MOOCs and technology to advance learning and learning research (Ubiquity symposium), Ubiquity 2014 (May) (2014) 2.

[13] P. Bours, S. Mondal, Performance evaluation of continuous authentication systems, IET Biometrics 4 (4) (2015) 220–226.

[14] L. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, Proceedings of the IEEE 77 (2) (1989) 257–286.

[15] D. Jurafsky, J. H. Martin, A. Kehler, K. Vander Linden, N. Ward, Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition, vol. 2, MIT Press, 2000.

[16] K. S. Killourhy, R. A. Maxion, Comparing anomaly-detection algorithms for keystroke dynamics, in: Dependable Systems & Networks, 2009. DSN'09. IEEE/IFIP International Conference on, IEEE, 125–134, 2009.

[17] N. Bakelman, J. V. Monaco, S.-H. Cha, C. C. Tappert, Keystroke Biometric Studies on Password and Numeric Keypad Input, in: European Intelligence and Security Informatics Conference (EISIC), IEEE, 204–207, 2013.

[18] M. J. Coakley, J. V. Monaco, C. C. Tappert, Keystroke Biometric Studies with Short Numeric Input on Smartphones, in: Biometrics Theory, Applications and Systems (BTAS) IEEE 8th International Conference on, IEEE, 2016.

[19] J. V. Monaco, N. Bakelman, S.-H. Cha, C. C. Tappert, Recent Advances in the Development of a Long-Text-Input Keystroke Biometric Authentication System for Arbitrary Text Input, in: European Intelligence and Security Informatics Conference (EISIC), IEEE, 60–66, 2013.

[20] M. Villani, C. Tappert, G. Ngo, J. Simone, H. S. Fort, S.-H. Cha, Keystroke biometric recognition studies on long-text input under ideal and application-oriented conditions, in: Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on, IEEE, 39–39, 2006.

[21] R. D. Malmgren, D. B. Stouffer, A. E. Motter, L. A. Amaral, A Poissonian explanation for heavy tails in e-mail communication, Proceedings of the National Academy of Sciences 105 (47) (2008) 18153–18158.

[22] C. C. Tappert, M. Villani, S.-H. Cha, Keystroke biometric identification and authentication on long-text input, Behavioral biometrics for human identification: Intelligent applications (2009) 342–367.

[23] G. Rupert Jr, et al., Simultaneous statistical inference, Springer Science & Business Media, 2012.

[24] P. Bours, Continuous keystroke dynamics: A different perspective towards biometric evaluation, Information Security Technical Report 17 (1) (2012) 36–43.

[25] T. Kobayashi, S. Haruyama, Partly-hidden Markov model and its application to gesture recognition, in: Acoustics, Speech, and Signal Processing, 1997. ICASSP-97., 1997 IEEE International Conference on, vol. 4, IEEE, 3081–3084, 1997.

[26] H. Ozkan, A. Akman, S. S. Kozat, A novel and robust parameter training approach for HMMs under noisy and partial access to states, Signal Processing 94 (2014) 490–497.

[27] S. O. Forchhammer, J. Rissanen, Partially hidden Markov models, IEEE Transactions on Information Theory 42 (4) (1996) 1253–1256.

[28] T. Iobayashi, J. Furuyama, K. Mas, Partly hidden Markov model and its application to speech recognition, in: Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on, vol. 1, IEEE, 121–124, 1999.

[29] S. Forchhammer, T. S. Rasmussen, Adaptive partially hidden Markov models with application to bilevel image coding, Image Processing, IEEE Transactions on 8 (11) (1999) 1516–1526.

[30] S. E. Levinson, L. R. Rabiner, M. M. Sondhi, An introduction to the application of the theory of probabilistic functions of a Markov process to automatic speech recognition, The Bell System Technical Journal 62 (4) (1983) 1035–1074.

[31] L. E. Baum, T. Petrie, G. Soules, N. Weiss, A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains, The annals of mathematical statistics 41 (1) (1970) 164–171.

Appendix A. Summary of POHMM parameters and variables

Table A.5: Summary of POHMM parameters and variables.

Parameter	Description
ψ,ω	Event types
i, j	Hidden states
$\mathbf{x}_1^N$	Observation sequence; $\mathbf{x}_n$ is the feature vector observed at time t_n
Ω_1^N	Event type sequence; Ω_n is the event type observed at time t_n
z_1^N	Sequence of hidden (unobserved) states; z_n is the hidden state at time t_n
M	Number of hidden states
m	Number of unique event types in Ω_1^N
$a\left[i,j\middle\vert\psi,\omega\right]$	Probability of transitioning from state i to j, given event types ψ while in state i and ω in state j
$\pi\left[j\middle\vert\omega\right]$	Probability of state j at time t_1, given event type ω
$\Pi\left[j\middle\vert\omega\right]$	Stationary probability of state j, given event type ω
$\mathbf{b}\left[j\middle\vert\omega\right]$	Emission distribution parameters of state j, given event type ω
$\gamma_n\left[j\middle\vert\omega\right]$	Probability of state j at time t_n, given event type ω
$\xi_n\left[i,j\middle\vert\psi,\omega\right]$	Probability of transitioning from state i at time t_n to state j at time t_{n+1}, given event types ψ and ω at times t_n and t_{n+1}, respectively

Appendix B. Proof of convergence

The proof of convergence follows that of Levinson et al. [30] which is based on Baum et al. [31]. Only the parts relevant to the POHMM are described. Let $Q\left(\theta,\dot{\theta}\right)$ be Baum's auxiliary function,

$$Q\left(\theta,\dot{\theta}\right)=\sum_{z_1^N\in Z}\ln u_{z_1^N}\ln v_{z_1^N} \tag{B.1}$$

where $u_{z_1^N}=P\left(\mathbf{x}_1^N,z_1^N\middle\vert\Omega_1^N,\theta\right)$, $v_Z=P\left(\mathbf{x}_1^N,z_1^N\middle\vert\Omega_1^N,\dot{\theta}\right)$, and Z is the set of all state sequences of length N. By Theorem 2.1 in Baum's proof [31], maximizing $Q\left(\theta,\dot{\theta}\right)$ leads to increased likelihood, unless at a critical point, in which case there is no change.

Using the POHMM parameters $\dot{\theta}$, $\ln v_{z_1^N}$ can be written as

$$\begin{aligned}
\ln v_{z_1^N} &= \ln P\left(z_1^N, \mathbf{x}_1^N | \Omega_1^N, \dot{\theta}\right) \\
&= \ln \dot{\pi}\left[z_1|\Omega_1\right] + \sum_{n=1}^{N-1} \ln \dot{a}\left[z_n, z_{n+1}|\Omega_n, \Omega_{n+1}\right] + \qquad \text{(B.2)} \\
&\quad \sum_{n=1}^{N} \ln f\left(\mathbf{x}_n; \dot{\mathbf{b}}\left[z_n|\Omega_n\right]\right) \qquad \text{(B.3)}
\end{aligned}$$

and similarly for $\ln u_{z_1^N}$. Then,

$$\begin{aligned}
Q\left(\theta, \dot{\theta}\right) &= \sum_{z_1^N \in Z} \Bigg\{ \ln \dot{\pi}\left[z_1|\Omega_1\right] + \sum_{n=1}^{N-1} \ln \dot{a}\left[z_n, z_{n+1}|\Omega_n, \Omega_{n+1}\right] \\
&\quad + \sum_{n=1}^{N} \ln f\left(\mathbf{x}_n; \dot{\mathbf{b}}\left[z_n|\Omega_n\right]\right) \Bigg\} P\left(z_1^N | \mathbf{x}_1^N, \Omega_1^N, \theta\right) \qquad \text{(B.4)}
\end{aligned}$$

and regrouping terms,

$$\begin{aligned}
Q\left(\theta, \dot{\theta}\right) &= \sum_{z_1 \in Z} \ln \dot{\pi}\left[z_1|\Omega_1\right] P\left(z_1 | \mathbf{x}_1^N, \Omega_1^N, \theta\right) \\
&\quad + \sum_{z_n^{n+1} \in Z} \sum_{n=1}^{N-1} \ln \dot{a}\left[z_n, z_{n+1}|\Omega_n, \Omega_{n+1}\right] P\left(z_n^{n+1} | \mathbf{x}_1^N, \Omega_1^N, \theta\right) \\
&\quad + \sum_{z_n \in Z} \sum_{n=1}^{N} \ln f\left(\mathbf{x}_n; \dot{\mathbf{b}}\left[z_n|\Omega_n\right]\right) P\left(z_n | \mathbf{x}_1^N, \Omega_1^N, \theta\right) \ . \qquad \text{(B.5)}
\end{aligned}$$

Finally, substituting in the model parameters and variables gives,

$$\begin{aligned}
Q\left(\theta, \dot{\theta}\right) &= \sum_{j=1}^{M} \gamma_1\left[j|\Omega_1\right] \ln \dot{\pi}\left[j|\Omega_1\right] \\
&\quad + \sum_{j=1}^{M} \sum_{i=1}^{M} \sum_{n=1}^{N-1} \xi_n\left[i, j|\Omega_n, \Omega_{n+1}\right] \ln \dot{a}\left[i, j|\Omega_n|\Omega_{n+1}\right]
\end{aligned}$$

$$+\sum_{j=1}^{M}\sum_{n=1}^{N}\gamma_n\left[j|\Omega_n\right]\ln f\left(\mathbf{x}_n;\dot{\mathbf{b}}\left[z_n|\Omega_n\right]\right) \tag{B.6}$$

The POHMM re-estimation formulae (Equations 13, 15, 16) follow directly from the optimization of each term in Equation B.6. Even when parameter smoothing is used, convergence is still guaranteed. This is due to the diminishing effect of the marginal for each parameter, $\lim_{N\to\infty}\tilde{\theta} = \theta$, where $\tilde{\theta}$ are the smoothed parameters.